

Serial Number 09/548,912**PATENT**
IBM Docket No. RAL9-00-0017**Amendments to the Specification:**

Amend page 2, paragraph beginning at line 20, as follows:

A1
Patent application S.N. 09/548,907 (~~tba - Docket RAL9-00-0010~~) filed concurrently by Brian M. Bass et al. and entitled "Method and System for Network Processor Scheduler". This patent is sometimes referred to herein as the Scheduler Structure Patent.

Amend page 2, paragraph beginning at line 23, as follows:

A2
Patent application S.N. 09/548,910 (~~tba - Docket RAL9-00-0014~~) filed concurrently by Brian M. Bass et al. and entitled "Method and System for Network Processor Scheduling Outputs Based on Multiple Calendars". This patent is sometimes referred to herein as the Calendar Scheduling Patent.

Amend page 3, paragraph beginning at line 3, as follows:

A3
Patent application S.N. 09/548,911 (~~tba - Docket RAL9-00-0015~~) filed concurrently by Brian M. Bass et al. and entitled "Method and System for Network Processor Scheduling Based on Calculation". This patent is sometimes referred to herein as the Calculation Patent.

Serial Number 09/548,912**PATENT****IBM Docket No. RAL9-00-0017**

Amend page 3, paragraph beginning at line 6, as follows:

A4
Patent application S.N. 09/834,141 (~~tba-Docket RAL9-00-0016~~) filed concurrently by Brian M. Bass et al. and entitled "Method and System for Network Processor Scheduling Based on Service Levels". This patent is sometimes referred to herein as the Service Level Patent.

Amend page 3, paragraph beginning at line 10, as follows:

A5
Patent application S.N. 09/548,913 (~~tba-Docket RAL9-00-0018~~) filed concurrently by Brian M. Bass et al. and entitled "Method and ~~System~~ System for Network Processor Scheduling Outputs using Disconnect/Reconnect Flow Queues". This patent is sometimes referred to herein as the Reconnection Patent.

Amend page 3, paragraph beginning at line 14, as follows:

A6
Patent application S.N. 09/546,651 (~~tba-Docket RAL9-00-0007~~) filed April 10, 2000 by Brian M. Bass et al. and entitled "Method and System for Minimizing Congestion in a Network". This patent is sometimes referred to herein as the Flow Control Patent.

Serial Number 09/548,912**PATENT**
IBM Docket No. RAL9-00-0017

Amend page 3, paragraph beginning at line 17, as follows:

A7
Patent application S.N. 09/547,280 (~~the Docket RAL9-00-0004~~) filed April 11, 2000
and entitled "Unified Method and System for Scheduling and Discarding Packets in
Computer Networks". This patent is sometimes referred to herein as the Packet Discard
Patent.

Amend page 8, paragraph beginning at line 19, as follows:

A8
The present invention overcomes the disadvantages and limitations of the prior art
systems by providing a simple, yet effective, way of handling information units or frames
coming out of a processing system and directing frames to output ports for dispatch to a
~~an~~ data transmission network. The present invention has particular application to a system
in which packets of variable length are being handled from a plurality of users and where
a level of service commitment has been made to at least some of the users.

Serial Number 09/548,912**PATENT**
IBM Docket No. RAL9-00-0017

Amend page 15, paragraph beginning at line 7, as follows:

A9
Weighted fair queuing (WFQ) calendars are used for so-called "best effort" service, and, when used in combination with a time-based calendar, a so-called "best effort peak" service. That is, a best effort service obtains no guaranteed bandwidth (where x bits of bandwidth is provided every unit of time), but competes with the other users for the bandwidth remaining after the guaranteed bandwidth customers are satisfied. This is a lower level of service than guaranteed bandwidth and normally significantly less expensive. The so-called "best effort peak" service provides a competitive sharing of excess bandwidth on a resource for requirements above the guaranteed level of service which a user has subscribed. So, a user might buy 10 Mbits per second of service along with a best effort for a peak of an additional 5 Mbits, for a total peak service of 15 Mbits of service (the last 5 Mbits of service are provided if and only if it is available and a fair sharing with other users if of the weighted fair queueing allows it).

Amend page 16, paragraph beginning at line 5, as follows:

A10
As described in the Scheduler Structure Patent a period (or clock interval) is defined as a scheduler_tick; this amounts to either 150 or 165 nanoseconds, depending on the response time of the hardware being accessed, but this could be greater or lesser

Serial Number 09/548,912**PATENT****IBM Docket No. RAL9-00-0017**

A10 Cont.

period depending on design parameters and hardware capabilities. During a scheduler_tick a flow queue is selected for service. The selection algorithm described results in the WFQ calendars being selected when all flow queues with minimum bandwidth specifications (sustained service rate) do not require service (i.e., the left over bandwidth managed by the scheduler is available for use by the WFQ calendars.) Put another way, the time-based calendars ~~queues~~ 220, 230 are served during each clock interval if they are scheduled and have information to transmit; if not, the WFQ queue 240 is served during the clock interval. In contrast to the timer based schedulers, where the total minimum bandwidth of the network processor may be managed, each WFQ calendar manages the best effort bandwidth for one target port. The purpose of the WFQ calendar ~~or queue~~ 240 is to share fairly amongst competing flow queues this best effort bandwidth. This is accomplished by accounting for the number of bytes transmitted by a flow queue that is selected for service and moving the flow queue within the calendar a distance from its present location based on the number of bytes transmitted. That is, the more bytes that a flow transmits during a scheduler_tick, the farther up the calendar (and more intervening flows and thus the longer the period) before the next service.

Serial Number 09/548,912**PATENT**
IBM Docket No. RAL9-00-0017

Amend page 17, paragraph beginning at line 15, as follows:

All
As shown in the flow chart of Fig. 5, during each scheduler_tick, the state of each target port queue is first examined. Each of the WFQ calendars is associated with a pair of ports; thus, WFQ Port 0 is associated with a higher priority port 0 and a lower priority port 0. If the target port queue's threshold has been exceeded on line 262, no further action is taken by that WFQ calendar during the scheduler_tick.. (This system provides a form of back pressure to limit the output, preventing frames from being sent out that the system can not handle.) If the target port queue's threshold has not been exceeded, the slot that is indicated by the current pointer is then examined. If the slot is found to be empty, then the current pointer may advance to the next non-empty slot to find a flow queue WFQ candidate. If all slots are found to be empty, the current pointer is unchanged and no candidate is found. If the slot is found to be non-empty within this one calendar, then the flow queue address in stored in the slot is the WFQ candidate for this port. Each of the 40 WFQ calendars will similarly be able to find a candidate for its associated target port queue.